

TITLE

"Procedure and system for the analysis and the evaluation of the conditions for accessing data communication networks, and relative computer program product"

DESCRIPTION

This invention refers to the techniques used to analyse and evaluate the conditions for accessing data communication networks such as the Internet.

To be precise, this invention has been developed with reference to its possible application to a service aimed at telecommunication networks for corporations, such as those commonly called "Corporate" networks or systems.

For a clear overall view of the criteria of organization and operation of this type of system see document WO-A-02/43406.

Given that the data network continues to emerge as a key element in the development of its own activities, corporate operators are expressing the need to be able to use several Internet Service Providers or ISPs to connect up to Internet, thus giving rise to a situation of "multi-homing". This is an alternative to the traditional solution, which uses a single provider and is defined "single-homing".

Two main factors lie behind this need: reliability and Internet connection performance. The use of two or more different providers makes it possible to increase Internet connection availability, and guarantee, for example,

greater potential for carrying out commercial transactions, or greater visibility on the outside world.

The use of various providers, for the same overall band, also improves the situation by appropriately balancing the traffic between the providers. For example, the decision to transit a certain type of traffic from/to a customer's site using provider A as opposed to provider B may result in an increase (or decrease) in performance depending on the provider's characteristics and routing policies.

In this type of application context, it is therefore a good idea to assess the opportuneness of changing from a "single-homed" situation to a "multi-homed" one, by using technical instruments and following objective criteria. In the event that a corporation decides to use a multi-homing connection architecture (connection to several providers), it is important to decide whether it would be a good idea to become an autonomous system (AS), and consequently implement the BGP (Border Gateway Protocol), or whether to use tools capable of handling the public addresses of several providers. The second solution assigns the addresses dynamically to corporate machines without incurring the cost of protocol management, which is costly both in economic terms and management terms as it requires high-level routers and highly qualified personnel.

The protocol called BGP (acronym for Border Gateway Protocol) is the tool currently used most to coordinate

routing between different Autonomous Systems (or AS) on the Internet. For a general discussion of the characteristics and methods of use of the BGP protocol, see the document entitled "A Border Gateway Protocol 4 (BGP-4) " by Y. Rekhter and T. Li, RFC 1771, T. J. Watson Research Center, Cisco, March 1995.

Document JP9181722 illustrates a system capable of creating the map of the autonomous systems (AS) that make up the Internet network. This is done by collecting the information from the router BGP tables.

Document US-A-6 243 754 illustrates a system for the dynamic selection of Internet providers. This system makes it possible to take appropriate measurements and dynamically select the provider that will provide the best performance at a given moment in time.

Document WO-A-02/17110 illustrates a solution that optimises the routing of traffic to a destination, when multiple routes are available. The relevant measurements are taken by analysing performance on the access routers, and then the system selects the best path for each destination and reconfigures the routing to the destination.

Finally, document WO-A-02/43322 illustrates a system that can be used if the network involved is part of a multi-homing configuration with various Internet providers. This system makes it possible to dynamically select the best link to the Internet each time or to balance the

traffic between the different links. This solution therefore presupposes that several Internet providers have already been selected.

The present invention aims to provide a solution that is capable of providing tools and information - of an objective nature - to evaluate the opportuneness of adopting a multi-homed architecture.

According to the present invention, this aim is reached thanks to a procedure having the characteristics specifically referred to in the annexed claims. The invention also refers to the relative system, as well as the corresponding computer program product that can be directly loaded into the internal memory of a numerical processor, and which includes parts of the software code required to implement the procedure as per the invention when the product is run on a processor.

In the preferred form of embodiment, the solution given in this invention includes two main stages.

The first stage traces the customer's Internet traffic to identify the main networks addressed by the traffic, the Internet sites most frequently visited, and the relative Autonomous Systems (AS) passed through. This can be done with hardware tools such as commercial probes or with suitable software agents on IP-level networking equipment, e.g. NetFlow marketed by the Cisco Corporation (USA).

The second stage traces the tree that represents the paths of the autonomous systems passed through by the

customer's traffic in order to decide whether (and with which provider) to connect up to the Internet in the multi-homed mode. This stage uses a tracing technique that requires the use of two modules.

The first module inputs the list of the most frequently visited Internet sites and then for each site it outputs the list of paths of the autonomous systems crossed to reach each destination. The second module aggregates all the information calculated by the first module and generates in output a tree representing all the paths of the autonomous systems crossed to reach all the destinations.

Three parameters should be indicated for each autonomous system: the percentage of use of the autonomous system, the average number of hops inside the autonomous system (AS) and the average amount of time spent inside the autonomous system.

The solution does not envisage the collection of information from the BGP table, nor the construction of the Internet network global map. It only envisages the construction of the tree of all the autonomous systems most frequently crossed by the traffic to all destinations. This in order to understand whether and with which providers the multi-homing Internet connection should be made.

Generally speaking, the solution given in the invention evaluates the need of a corporation to use several providers to access the Internet, thus avoiding the

necessity of having to dynamically choose the best provider. All the destinations, in fact, are considered globally in order to decide not the best path but whether it is to the corporation's advantage to have several paths represented by as many Internet providers. The best path or paths will be selected subsequently according to criteria chosen by the user. The advantage of having several links with the Internet will then be objectively evaluated, and then, if necessary, the providers to be used to make the links identified.

In the preferred form of embodiment, the solution given in the invention provides two macro-categories of essential information for the decisional process:

- tracing of the customer's Internet traffic, which makes it possible to identify the main networks the traffic addresses (and the relative autonomous systems) and the relative volume;
- tracing of the sequence of the autonomous systems crossed by the customer's traffic in order to decide whether (and with which provider) to activate an Internet connection.

A series of measurements, taken with a probe or with router functions referred to earlier, makes it possible to obtain the customer traffic matrix and subsequently process the information to identify the target autonomous systems and the relative paths. All this constitutes the objective base of the decisional stage.

The invention will now be described here below by way of example, and not of limitation, with reference to the attached drawings in which:

- figures 1 and 2 illustrate the reference scenarios of a corporation that visits the Internet with a single-homed approach and a multi-homed approach respectively,

- figures 3 to 5 represent, in the form of "cake" diagrams, lists of the most frequently visited networks, most frequently sites and the main destination autonomous systems respectively,

- figure 6 illustrates the corresponding paths of the autonomous systems crossed by the customer's traffic,

- figure 7 illustrates the corresponding performance values,

- figure 8 is a flow diagram showing the development of the procedure according to the invention, and

- figure 9 illustrates a possible example of a table generated by a traceroute function during the implementation of the solution as given in the invention.

Figures 1 and 2 illustrate the reference scenario of a corporation (herein represented by its local network or LAN) in relation to the Internet access made through a single provider (ISP#1 in figure 1) and through various providers (providers IPS#1, ISP#2 and ISP#3 in figure 2). These are therefore the scenarios commonly called "single homed" and "multi-homed".

A corporation that currently uses a single-homed configuration and wishes to have a second Internet access through another provider must answer a certain number of questions when it starts assessing the need to move on to a multi-homed scenario.

In particular, it is important to be aware of the following when thinking about changing from a single-homed scenario to a multi-homed scenario: how the corporation's traffic is distributed, especially as regards the networks towards which the greatest volume of traffic is directed; which autonomous systems are crossed by the traffic, and in particular the autonomous systems in which the traffic terminates; who are the main visitors to the corporate web/e-commerce sites; and from which autonomous systems (AS) they originate.

It is especially important to identify which providers should be used to make new connections to the Internet when selecting a multi-homed scenario.

The solution described herein not only supplies the information on the requirements and most frequently used main traffic lines, necessary to make a decision on whether to change to multi-homing access but also, in the event that multi-homing has already been implemented, it makes it possible to define alternative connection and routing policies with various providers, and if necessary helps decide whether to change one or more providers or not.



The solution given in this invention aims to obtain the following for both scenarios described above:

- traffic measurements, such as band use measurements, traffic volume, congestion levels, load balancing, and indications on the most frequently visited networks;

- a list of autonomous systems (AS) most frequently crossed by the corporation's local LAN network to the Internet;

- percentages of use of the various autonomous systems towards the Internet, and

- statistics to analyse who are the main visitors to the corporation's local sites and from which autonomous systems these visits originate.

To do this, the solution in this invention uses various analysis tools.

These may be for example probes - of the type normally on sale - that can be used to obtain measurements and traffic statistics (most visited networks, traffic volume, congestion levels, use of links). Alternatively, the solution in the invention can use software agents on NetWorking IP equipment, e.g. NetFlow<sup>TM</sup>, which has been mentioned earlier.

The information obtained is then processed so as to trace the paths of the autonomous systems most frequently visited, and to determine which providers are most involved, by analysing the percentage of use of the autonomous systems.

The examples shown in figures 3, 4 and 5 illustrate various diagrams, which can be obtained as shown later, and which show how the incoming/outgoing traffic to/from the LAN network examined is subdivided. They give the information relating to the destination networks (figure 3), to the percentage of traffic involved (figure 4), and to the pertinent autonomous system (figure 5).

As shown in figure 6, a tree can be built with leaves that are the subnetworks that are the destinations of the traffic of the LAN involved. The corresponding report illustrated in figure 7, shows the autonomous systems crossed to reach the various subnetworks and gives information on how the traffic is divided (e.g. in percentage) at different levels of the tree.

The information in figure 6 helps choose which providers can be used to implement multi-homing policies or (in the event that a multi-homing scenario has already been implemented), to change the Internet connections already active.

Once the list of autonomous systems crossed by the customer's traffic has been drawn up, the average amount of time spent and the average number of hops inside the autonomous system can be found for each one as shown in figure 7.

Using the information described above, and proceeding as illustrated below with reference to figure 8, a report

can be generated for the final user containing the following information:

- tracing of the customer's traffic to the Internet with identification of the main networks with which the traffic is involved (and the relevant autonomous systems), as well as the relative volume, and

- tracing of the sequence of autonomous systems crossed so as to determine whether and with which providers the Internet links should be made.

Stage A of figure 1 is aims at tracing the Internet traffic of the customer's LAN network and basically includes a step, marked with A1 in figure 8, that monitors the Internet access links for the collection of traffic data. The results, collectively referred to with 100, correspond to the list of IP networks and addresses most frequently visited by the customer's Internet traffic.

The subsequent stage, referred to with the letter B, includes the evaluation of the paths of the autonomous systems crossed (AS path) by the customer's traffic. The first step towards this is indicated with B1 and traces the autonomous systems crossed a sufficiently high number of times for each destination network/address in the list marked with 100.

The result of this, indicated with 102, is the list of paths of the autonomous systems crossed to reach each destination.

The next step, indicated with B2, then aggregates all the information collected. This processing firstly generates a group of results, 103, which corresponds to a unique tree made up of the paths of the autonomous systems crossed by the customer's traffic, indicating the subdivision, in percentage, of the traffic on each path.

A second set of results, 104, is a table showing the calculation of the average number of hops inside each autonomous system and the calculation of the average amount of time spent inside each autonomous system.

During stage A in figure 8, which identifies the IP networks that generate the most traffic from/to the network under examination, the solution employs systems of the type used to monitor the use of the links, trace the customer's traffic, and identify the main traffic lines, the most frequently visited Internet sites, the most frequently used protocols, and the busiest times of day.

To do this, the solution employs specific, known hardware probes able to provide information on the band use of an individual link, on the volume of data, on the subdivision according to protocol, IP address, and the traffic matrix between the network under examination and the Internet network. This makes it possible to identify which Internet sites are most frequently visited by the customer network, and consequently which are the main networks addressed by the customer's traffic. The incoming traffic is also taken into account, which gives information

on the origin of those who connect up to the customer's network (server web, server ftp, etc.). These products generate a report including the list of IP addresses most frequently visited, and constitute the group of input data to be used for the subsequent stage of analysis and post-processing.

Alternatively, as has already been mentioned, software agents, such as NetFlow, operating on the Internet access routers, can be used. These software agents can be used to trace the incoming/outgoing traffic to/from the customer's router interface that connects to the Internet, and to identify the main traffic lines. All this can be done by analysing the operating status of the router in terms of CPU load and available memory. If this solution is adopted, it is necessary to decide where to export the statistics autonomously created by the router, and identify a machine onto which these data can be imported.

Stage B in figure 8 is used to obtain the information relating to the autonomous systems crossed by the customer's traffic to reach the destination addresses.

As already mentioned, this involves performing steps B1 and B2, and using an autonomous system tracing system basically consisting of two modules.

The first module inputs file 100 containing the IP addresses representing the sites most frequently visited by the customer from stage A. It sends a traceroute ICMP message (Internet Control Message Protocol) several times

to each destination site (with a configurable frequency), and each time it traces the path to reach this destination.

The path in question is expressed as a sequence of IP addresses. Figure 9 gives an example of the data table generated by this traceroute function.

In order to relate the aforesaid IP addresses to the corresponding autonomous systems (AS), software script is used to interface with databases like the ones handled by RIPE (Réseau IP Européen), ARIN (American Registry for Internet Numbers) and APNIC (Asia Pacific Network Information Center), i.e. by the three organisations that supervise the handling of problems regarding the Internet at a European, American and Asia-Pacific level.

The second module aggregates all the information calculated by the first module, generates a unique tree of autonomous system paths crossed by the customer's traffic to reach all the destinations, and gives three parameters for each autonomous system, i.e. percentage of use of the autonomous system, average number of hops inside the autonomous system, average amount of time spent inside the autonomous system.

To return to the methods of tracing the autonomous systems in greater detail, the first module, as mentioned previously, performs the following operations:

- inputs a list of destination URL (or host IP addresses or network IP addresses): an input file can be

hypothesised with a simple list of URL separated by a separator, e.g. one in each row;

- performs the traceroute function several times according to a configurable frequency (e.g. every 30 minutes) of the path to each element (URL, host address or IP address) in the list;

- invokes a remote identification service (whois?) of the aforementioned databases RIPE, ARIN, APNIC, for each IP address generated by the afore-said traceroute function, in order to obtain the name of the autonomous system to which the IP address belongs, and the number of the AS to which the IP address belongs; and

- organises the data obtained into data output format.

The format in question generally envisages output files for each destination IP address, in which each file is a list of lines or rows with identical structure.

Each line contains the path of the AS crossed to reach a single destination, and is obtained by a single traceroute command used on the destination address. Each output file contains as many lines as traceroutes performed according to a configurable frequency and each line is an ordered sequence of elements separated by a separator such as ";" (semi-colon).

Each element represents the data relating to an autonomous system of the path. In the preferred embodiment, the format of each element includes the following in the order given:

- the order number of the autonomous system following the IP address sequence supplied by the traceroute function,

- the text name of the autonomous system,

- the identification number of the autonomous system,

- the number of hops that the single tracing command has measured inside the autonomous system (several IP addresses may belong to the same AS), and

- the time interval spent in the autonomous system, normally expressed in milliseconds, measured by the single tracing command.

Two typical examples of input and output files of the module under examination are given below.

Example of an Input File:

www.cisco.com

www.telecomitalia.it

193.206.129.254

193.206.132.146

193.206.132.178

162.40.1030.0

Example of an Output File:

- 1,AS\_alfa,AS100,3hop,0.326msec;2,AS\_beta,AS160,7 hop,  
0.36 msec;3,AS\_gamma,AS200,2 hop, 0.776 msec;
- 1,AS-alfa,AS100,3 hop, 0.326 msec 1;2,AS\_epsilon,  
AS180 4 hop, 1.3 msec;
- 1,AS\_alfa,AS100 ,3 hop ,0.526 msec ;2,AS\_beta, AS160,7  
hop, 0.38 msec;3,AS\_epsilon,AS180 4 hop, 1.3 msec.



The module uses the *whois* remote service of the databases RIPE, ARIN, APNIC for each IP address documented by the traceroute function to obtain information relating to the name and number of the autonomous system in question. All the other information (i.e. the number of hops inside each autonomous system and the number of milliseconds spent in each autonomous system) is processed by the module starting with an analysis of the output of each tracing operation.

Figure 9 gives an example of output of the aforesaid traceroute command.

Once it has been determined which autonomous system each hop belongs to by means of the information from the *whois* service, it is easy to calculate the average time (approximate) for the transit of packets in the autonomous system and the number of internal hops.

The methods that can be used by the module to ascertain the input of a list of destination hosts (URL) and to perform the traceroute function for each one of them sequentially, can be improved in terms of rapidity by generating a certain number of processes to each of which a traceroute command can be given in parallel. The original set of destinations can be divided and a subset of destinations attributed to each of the processes generated.

This results in the IP addresses table, and consequently the information on the autonomous systems, being obtained more quickly. Generally speaking, the

execution times are approximately and on average inversely proportional to the number of parallel processes started, at least until the number is not equal to the original number of URL. At each module run, it is also possible to dynamically give a value of the parallel processes in relation to the number of input URL's, by making this number vary from one to the original number of URL's.

It should also be appreciated that it is not generally necessary to access the *whois* remote service for each IP generated by the *traceroute* function. Bearing in mind that during these interrogations it is extremely probable that the first hops in a path already travelled will appear to be revisited, it is clear that it is a good idea to create and use a local cache memory that can store the correspondence between the IP addresses and the information relating to their autonomous systems. This means that the *whois* remote service interrogations do not need to be carried out again, if the last interrogation took place only a short time before.

Given that the information in the external databases, such as the RIPE, ARIN and APNIC databases, RIPE, ARIN e APNIC), may vary, once this information has been entered (and become redundant) inside the cache memory, it cannot be considered definite. A cache memory information update function must therefore be included.

At a configurable frequency, this function indicates for how long the information has not been updated and, for

information considered no longer valid because not updated for a long time, it interrogates the external databases and updates the information in the cache memory.

There may be cases when these databases have no information on the autonomous system relating to a given IP address.

This information can be obtained with other tools, e.g. by consulting web sites, and the data not obtainable from interrogating the aforesaid databases can be added to the local cache.

The second module referred to previously inputs one or more text files generated by the first module, and its objective is to aggregate the autonomous system (AS) paths for all the destinations.

Processing then traces the aggregated paths for all the destinations. It then outputs a tree with leaves that are the destination subnetworks of the customer's traffic and branches that are the autonomous systems crossed by the traffic. This representation highlights the autonomous systems crossed to reach the various subnetworks and shows how the traffic is divided (in percentage) at the different tree levels, in the terms shown in figure 6.

This second module therefore has the following aims:

- represent the aggregated path list (AS path),
- calculate the path crossing percentage to all the URL's obtained by aggregating the information received from the first module,

- generate in output a legible text format,
- generate in output a format that can be integrated with external tools or software,
- generate in output a table including the calculation of the average number of hops inside each autonomous system and the calculation of the average time spent inside each autonomous system.

This second module inputs and processes one or more text files generated by the first module seen previously. To satisfy its aim of constructing a tree with leaves representing the destination subnetworks of customer traffic with indications of the autonomous systems crossed to reach these subnetworks, and to provide information on how the traffic is divided (in percentage) at different tree levels, the first step for this second module is to generate a data structure representing the paths generated by the first module in the central memory.

In its preferred embodiment, the representation used is an aggregated list (LA), or a group of prefix-aggregated lists. An aggregated list represents a variable number of lists (of variable length), of nodes (autonomous systems, in this particular case) that share the common maximum prefix.

For example the following lists can be considered:

- abcdefghi
- abcde123
- ab123

These lists can be represented as follows with an LA:

a-b+c-d-e+f-g-h-i

| +1-2-3

+1-2-3

The example shows that the nodes <1,2,3> appear twice in the LA.

Therefore, if the first module generates the following output (where for the sake of simplicity the information about the number of hops and time inside each AS is not given):

1, AS-ISP1, ASnumber1	2, XANGE-NET ASnumber3	3, AS-ISP3, ASnumber7	4, AS-US-ISP ASnumber9
1, AS-ISP1, ASnumber1	2, WEB-NET, ASnumber4		
1, AS-ISP1, ASnumber1	2, XANGE-NET, ASnumber3	3: AS-GlobalISP, ASnumber8	4, AS-EDU-net, ASnumber10
1: AS-ISP2, ASnumber2	2 new-NET ASnumber5		
1, AS-ISP2, ASnumber2	2, Other-NET, ASnumber6		

then the second module must build up the following aggregated lists

ASP-ISP1, ASnumber1	XANGE-NET, ASnumber3 (40%)	AS-ISP3,	AS-US-ISP,
---------------------	----------------------------	----------	------------

(60%)		Asnumber7 (20%)	ASnumber9 (20%)
	+	AS-GlobalISP, Asnumber8 (20%)	AS-EDU-net, ASnumber10 (20%)
+	WEB-NET-WEB, Asnumber4 (20%)		
AS-ISP2, ASnumber2 (40%)	new-NET, ASnumber5 (20%)		
+	Other-NET, ASnumber6 (20%)		

The percentage of traffic indicated next to each autonomous system represents, in terms of overall traffic, the percentage of traffic that passes through the autonomous system. For example, starting from the output of the first module, it is possible to deduce that since there are 3 examples of AS-ISP1 and AS-ISP1 at first level in the period of time analysed, 60% of the total traffic transited on the first autonomous system and 40% on the second.

In the 60% of the traffic generated in AS-ISP1, since there are 2 examples of XANGE-NET at the second level with prefix AS-ISP1 and one example of web-NET with the same prefix, we can deduce that 40% of this traffic transited to XANGE-NET and the remaining 20% to web-NET.

Similar considerations can be made for the levels. In this way it is possible to know how the customer's Internet traffic is divided between the various autonomous systems

in order to choose, if necessary, the provider with which to implement a multi-homed configuration, or if the corporation has already adopted a configuration of this type, to decide whether to confirm the agreements with the current providers or whether to use other providers.

In addition to the first output, the second module generates a summary table, starting from the input of the first module, containing the list of all the autonomous systems analysed.

For each one of these, the average number of hops and the average amount of time spent inside each autonomous system is calculated.

An example of a table of this type is given below:

AS Name	ASnumber	Time	Number hops
XANGE-NET	ASnumber3,	22.66 ms	1.02
AS-ISP1	ASnumber1,	55.75 ms	5.88
AS-GlobalISP	ASnumber8,	65.42 ms	4.17
AS-ISP3	ASnumber7,	15.96 ms	4.88
AS-US-ISP	ASnumber9,	16.89 ms	2.50
AS-ISP2	ASnumber2,	96.65 ms	1.61
WEB-NET	ASnumber4,	0.00 ms	1.00
New-NET	ASnumber5,	58.40 ms	1.00
AS-EDU-net	ASnumber10,	48.20 ms	1.20
Other-NET	ASnumber6,	13.2 ms	2.20

Naturally, while keeping to the principal of the invention, the details regarding the construction and the embodiments of the invention may vary considerably with respect to what has been described and illustrated, without however departing from the scope of the present invention.



**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☒ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**